

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Consolidation of Allocated Memory to Reduce Power
Consumption**

Inventor(s):
Steven C. Woo
Pradeep Batra

CROSS-REFERENCES TO RELATED APPLICATIONS

This U.S. Nonprovisional Patent Application is a continuation of Application No. 09/919,373, filed on July 30, 2001.

TECHNICAL FIELD

This invention relates to power conservation in memory devices and systems.

BACKGROUND

Dynamically refreshed memory, usually referred to as dynamic random access memory or DRAM, is a type of memory device found in many different computing devices. A typical DRAM device may have millions, billions or even more DRAM memory cells. A DRAM memory cell is commonly formed by a single transistor and an associated capacitance. The capacitance is charged to a voltage that indicates a bit value of either "0" or "1". The capacitance loses its charge rather quickly, bringing about the need for periodic refreshing.

In many computer systems, the power consumption of DRAM memory is insignificant compared to other system components such as hard disks, high-performance microprocessors, active matrix displays, CRTs, etc. However, in other computer systems, such as the newly emerging and evolving class of mobile devices known as "handhelds" or "PDAs" ("personal digital assistants"), the power consumption of the DRAM memory is significant as compared to other components in the computer system. In comparison to many of the more traditional types of computers, such as desktop or personal computers, many mobile computing devices, are smaller, less capable, and use components that

1 consume less power. For example, many of these systems have small,
2 monochromatic displays, low performance CPUs, and no hard disks. Some of these
3 mobile systems, furthermore, rely on batteries for their operating power. As a
4 result of these factors, power consumption of memory subsystems has become
5 more of an issue in these devices; there is a strong need to reduce memory power
6 consumption and to thereby extend the time between required battery replacement
7 or recharging.

8 Memory devices with power management features are becoming available
9 to address this need. For example, DRAMs are available that support various
10 different reduced power modes. However, power savings come at the cost of
11 performance. Typically, a greater penalty in access speed is imposed at each
12 increasing degree of power savings. Thus, decisions regarding whether to invoke
13 power-saving features in a DRAM should be made intelligently. Typically, it is
14 desired to initiate a low power mode in a particular memory device only when that
15 memory device is not currently in use and is not anticipated to be in use in the near
16 future.

17 It is difficult, however, to anticipate the future need for accessing any
18 particular region of memory. Furthermore, modern operating systems typically
19 allocate memory without regard to memory device boundaries, making it difficult
20 to find a single memory device that can appropriately be set to a reduced power
21 mode without significantly impacting overall system performance. More
22 specifically, typical memory allocation schemes often result in a highly
23 fragmented memory space, with allocated pages of memory spread more or less
24 randomly across the available range of physical memory. Because allocated
25 memory is normally spread across all of the available devices, none of the devices

1 can be put into a reduced power mode without seriously impacting memory
2 performance.

3 An article entitled "Power Aware Page Allocation," by authors Alvin R.
4 Lebeck, Xiaobo Fan, Heng Zeng, and Carla Ellis, in Proceedings of the Ninth
5 International Conference on Architectural Support for Programming Languages
6 and Operating Systems (ASPLOS-IX) (November, 2000), discusses the
7 desirability of clustering memory page allocations into the minimum number of
8 memory devices. Such clustering allows other devices to be put into reduced
9 power modes.

10 Described below are specific techniques for minimizing the number of
11 actual DRAM devices being used in a system at any particular time. Such
12 techniques can be implemented in existing systems with very little overhead, while
13 potentially achieving significant power savings.

14 15 **BRIEF DESCRIPTION OF THE DRAWINGS**

16 Fig. 1 is a block diagram of a memory system that incorporates aspects of
17 the invention.

18 Fig. 2 shows a logical-to-physical mapping table used in conjunction with
19 the memory system of Fig. 1.

20 Fig. 3 is a block diagram of a memory device that can be utilized in
21 conjunction with the memory system of Fig. 1.

22 Fig. 4 is a block diagram of an alternative memory system that incorporates
23 aspects of the invention.

24 Fig. 5 shows a virtual-to-physical mapping table used in conjunction with
25 the memory system of Fig. 4.

1 Fig. 6 illustrates a free page list used in conjunction with the memory
2 system of Fig. 4.

3 4 **DETAILED DESCRIPTION**

5 **Memory Consolidation by Logical-to-Physical Re-mapping**

6 Fig. 1 shows pertinent portions of a computer system 10, including a CPU
7 12, a memory controller 14, and memory devices 16. Although the memory
8 controller and memory devices are shown to be separate entities in this figure, the
9 same techniques apply for memory controllers that are integrated into the CPU, as
10 well as memory that is integrated with either the controller and/or the CPU.

11 The computer system also includes an operating system 18 and one or more
12 applications or application programs 20. The operating system and applications
13 are typically initially stored on some form of non-volatile memory (not shown).
14 They are subsequently loaded into executable memory and executed by CPU 12.
15 Devices 16 form at least part of the executable memory. In many cases, the
16 computer system implements a virtual memory system, so that only portions of the
17 operating system and applications are actually present in physical memory at any
18 given time.

19 The architecture of Fig. 1 is typical of many computers and computer-like
20 devices, and is not limited to conventional desktop systems or even to
21 conventional portable computer systems. Many types of devices, such as
22 entertainment and game devices, industrial control devices, and others either use
23 an architecture such as this or can be easily adapted to use such an architecture.

24 The operating system is typically an off-the-shelf, general-purpose
25 operating system that provides low-level management functions and support for

1 higher-level application programs. However, the operating system might
2 alternatively be a custom application or program designed for a particular,
3 specialized purpose, and might itself perform the specialized functions that would
4 in other cases be performed by separate application programs.

5 In the described embodiment, memory devices 16 have dynamically
6 refreshable memory cells. Such devices are typically referred to as DRAMs
7 (dynamic random access memory), or DRAM devices. Other types of memory
8 devices can, however, also benefit from the techniques described herein.

9 Memory devices 16 can be arranged in many different configurations.
10 Typically, memory devices are arranged in sets or ranks, with each device storing
11 a single bit or other portion of a digital word. The bits or portions are then
12 combined or grouped when writing to or reading from a particular memory
13 address. Decisions regarding whether devices are to be operated in reduced power
14 modes are typically made with respect to all of the devices of a particular set or
15 rank. For purposes of discussion, it is assumed that each block 16 in Fig. 1
16 represents a set or rank of one or more memory devices, which are configured as a
17 group in different power modes.

18 Memory controller 14 acts as an interface between CPU 12 and memory
19 devices 16. Memory controller 14 has refresh logic 21 that is configured to
20 periodically refresh the memory cells of the memory devices. Although not
21 shown, each of memory devices 16 has multiple dynamically refreshable memory
22 cells, arranged in rows. In operation, memory controller 14 can receive memory
23 instructions from various sources, including but not limited to, operating system
24 18, CPU 12, a graphics adapter (not shown), and/or other sources. Memory
25 controller 14 responds to the instructions by performing various memory

1 operations such as, for example, reads and writes. When performing a memory
2 operation, the memory controller specifies particular memory cells by means of a
3 physical address or a combination of addresses (such as row and column
4 addresses). For purposes of the following discussion, these memory addresses
5 will be referred to as physical addresses, and the available range of physical
6 memory addresses will be referred to as the physical address space.

7 Instructions received by memory controller 14 from CPU 12 specify
8 addresses in terms of an address space that is normally referred to as a physical
9 address space, but which for purposes of the following discussion will be referred
10 to as a *logical* address space. The term “physical address space” will be used to
11 refer to the range of addresses occupied by memory devices 16. In many prior art
12 systems, these two address spaces are equivalent. In the described embodiment,
13 however, there is a variable mapping between logical addresses used between
14 CPU 12 and memory controller 14 and physical addresses used between memory
15 controller 14 and memory devices 16.

16 The logical address space is a linear range of addresses that is mapped by
17 memory controller 12 to the physical address space occupied by memory devices
18 16. As will be seen, mappings of the logical address space to the physical address
19 space are variable—they change over time. However, this variability is
20 transparent to the CPU. To the CPU, the logical address space appears simply as a
21 linear address space. All memory references by the CPU are specified in terms of
22 the logical addresses that form the logical address space.

23 More specifically, memory controller 14 has mapping logic or address
24 translation logic 22 that maps or translates from logical addresses specified in
25 memory instructions to physical addresses used to access physical memory

1 devices 16. Such mapping is performed in a manner that reduces or minimizes the
2 number of physical memory devices or ranks of memory devices that are in use at
3 any given time. After minimizing the number of physical memory devices that are
4 in use, the memory controller sets other memory devices—those that are either
5 unused or are subject to relatively less usage—to reduced power modes.

6 More specifically, memory controller 14 repeatedly performs actions of (a)
7 identifying portions of the logical address space based on usage, (b) mapping the
8 identified portions of the logical address space to physical memory in a manner
9 that reduces the number of physical memory devices referenced by the identified
10 portions of the logical address space, (c) identifying one or more memory devices
11 that are not referenced by the identified portions of the logical address space, and
12 (d) setting said one or more identified memory devices to one or more reduced
13 power modes. These actions will be described in more detail in the following
14 discussion.

15 Fig. 2 shows a logical-to-physical address translation table 30 such as might
16 be used to implement address translation logic 22. Logical-to-physical address
17 translation can be done on at the byte level (that is, each byte of the logical address
18 space can be mapped to a different byte in the physical address space), but in
19 practice this requires the translation table to be too large. In practice (and in the
20 preferred embodiment), entries of the translation table correspond to groups of
21 bytes in the logical address space (for example, the number of bytes in a physical
22 page, a DRAM page, a DRAM bank, or even a rank of DRAMs). Throughout the
23 remainder of this description, the term “logical page” will be used to represent
24 groups of addresses in the logical address space that map to groups of addresses in
25 the physical address space.

1 Table 30 has a plurality of mappings or entries 31, each of which indicates
2 a logical page address or number in a first column 32 and a corresponding targeted
3 physical page address or number in a second column 33. For purposes of
4 discussion, a particular physical page address or number will be said to be the
5 “target” of a logical-to-physical mapping whenever there is an entry in table 30
6 that maps to said particular physical page address or number.

7 When a received memory instruction specifies a particular logical address
8 or logical page number, memory controller 14 translates it to a physical address or
9 physical page number by referencing the appropriate entry of the address
10 translation table. In accordance with the techniques described herein, memory
11 controller 14 repeatedly and/or periodically modifies the mappings of table 30 to
12 reduce the number of physical memory devices that are targeted by the mappings.

13 14 **Monitoring Memory Instructions to Identify In-Use Memory**

15 In one embodiment, memory controller 14 identifies highly used portions
16 of the logical address space by monitoring memory instructions and keeping track
17 of which logical addresses or logical memory pages are specified most frequently
18 and/or most recently in the memory instructions. In this embodiment, memory
19 controller 14 is configured to periodically re-map the logical address space to
20 physical memory, to reduce the number of physical memory devices referenced by
21 the identified, highly-used portions of the logical address space. Specifically, the
22 address translation entries are periodically re-calculated so that the most frequently
23 and/or recently used logical memory addresses map to physical memory that is
24 located in the fewest possible number of physical memory devices or ranks.
25 Depending on system constraints which set the maximum power consumption of

1 the memory system, the most frequently and/or recently used logical addresses can
2 be mapped to devices that consume more power (and which have the highest
3 performance), such as devices operating in “Attention” and “Standby” modes
4 available in “Direct RDRAM” memory devices manufactured by Rambus Inc., of
5 Los Altos, California. Logical addresses that are used less frequently and/or less
6 recently can be mapped to devices that consume less power (and which have lower
7 performance), such as devices operating in “Nap” or “Powerdown” modes of
8 “Direct RDRAM” memory devices. In conjunction with this re-mapping process,
9 memory content is copied or moved as appropriate so that all logical memory
10 addresses will continue to reference the same data even though the data might now
11 be in a different physical location.

12 13 **Operating System Notifications of In-Use Memory**

14 In other embodiments, memory controller 14 receives explicit notifications
15 from operating system 18 regarding allocations and de-allocations of memory.
16 These notifications preferably indicate one or more logical memory pages that are
17 being allocated or de-allocated.

18 Typically, an operating system includes facilities for dynamically allocating
19 and de-allocating memory. When loading an application, for example, an
20 operating system typically designates specific areas of memory for the code of the
21 application and specific areas of memory for use by the program in storing data.
22 Allocation and de-allocation typically involve maintaining one or more tables or
23 other data structures indicating those areas of memory that have been designated
24 for use in this manner. Such areas are typically identified within such tables or
25 data structures by their memory addresses—by their physical memory addresses in

1 most prior art systems, but by their *logical* addresses in the embodiments
2 described herein.

3 Memory allocation can also take place as a result of an application program
4 requesting the use of additional memory during actual execution of the application
5 program. In response to requests such as this, the operating system designates
6 areas of memory (which in the described embodiment comprise *logical* memory)
7 for exclusive use by the requesting application programs.

8 In operating systems that support virtual memory, allocation of physical
9 memory typically takes place at a lower level. Systems such as this create an
10 individual *virtual* address space for each of multiple application programs. Each
11 virtual address space is very large—typically much larger than the amount of
12 available physical memory.

13 In systems such as this, the operating system typically allocates virtual
14 memory to requesting application programs. When such virtual memory is
15 allocated, the operating system creates a translation entry or “mapping” between
16 an allocated range of virtual memory addresses and a corresponding range of
17 physical memory addresses—in the embodiments described herein, each mapping
18 is between a range of virtual memory addresses and a corresponding range of
19 *logical* memory addresses. Each translation entry or mapping translates from a
20 virtual or source address to a logical or target address. The operating system
21 maintains a translation or mapping table that contains all current translations or
22 mappings.

23 When an application program subsequently references a virtual memory
24 address, the operating system and CPU use the translation table to translate from
25 the virtual address to the logical address, and the actual memory access is made to

1 the indicated logical address. This translation process is transparent to the
2 application program.

3 In order to make each virtual address space appear relatively unlimited, the
4 operating system makes use of a mass storage medium such as a hard disk, which
5 is typically referred to as secondary storage or secondary memory to distinguish it
6 from primary or physical memory. Secondary storage is usually relatively slow to
7 access, but normally has a capacity much larger than that of primary memory. The
8 operating system monitors memory usage and when portions of virtual memory
9 are not being used, the data from the corresponding portions of physical memory
10 is moved to secondary storage. Thus, at any given time, some portions of virtual
11 memory will correspond to portions of physical memory, and some virtual
12 memory will correspond to portions of secondary memory.

13 If an application program attempts to access a portion of virtual memory
14 that is currently held in secondary storage, there will be no appropriate entry in the
15 translation table. This is referred to as a "miss," in response to which the
16 operating system intervenes, loads the appropriate data back into physical memory
17 and creates an appropriate translation entry in the translation table. After this is
18 accomplished, the control is returned to the application program, which accesses
19 the memory in its normal fashion.

20 The process of moving data between primary and secondary storage is
21 referred to as memory "swapping" and normally takes place on an ongoing basis.
22 As part of this process, the operating system maintains and updates its virtual-to-
23 logical address mappings so that any reference to a virtual memory address will be
24 translated to the appropriate logical address. The virtual-to-logical mappings
25 change frequently in response to memory swapping.

1 Thus, in systems that support virtual memory, the operating system
2 allocates *virtual* memory to requesting application programs. Prior to use,
3 however, the operating system loads needed portions of the virtual memory into
4 portions of physical memory, and provides address translations between virtual
5 and logical memory addresses. In systems such as these, logical memory can be
6 considered to be allocated whenever it is the target of an active virtual-to-logical
7 memory mapping as described above. The operating system is configured to
8 notify controller 14 when portions of the logical address space becomes allocated
9 in this fashion.

10 Regardless of the method of memory allocation, the operating system is
11 configured to identify allocated portions of the logical memory space to controller
12 14. Specifically, the operating system informs the memory subsystem of the
13 specific addresses or address ranges that are being allocated or de-allocated.
14 These addresses are specified in terms of the logical address space described
15 above, although such addresses would normally be thought of as “physical”
16 addresses from the point of view of the operating system.

17 In response to receiving a notification of a page allocation, memory
18 controller 14 creates a new mapping or translation entry 31 for use by address
19 translation logic 22. As new mappings are created, they are created in a manner
20 that tends to reduce the number of additional memory devices or ranks targeted by
21 the new mappings. Thus, if possible, the new entry is created so that it targets
22 physical memory from a memory device or rank that is already targeted by an
23 existing translation entry. This tends to consolidate physical memory usage in the
24 fewest possible memory devices, allowing non-targeted memory devices to be set
25 to reduced power consumption modes. If it is not possible to allocate the new

1 page from an already targeted memory device or memory rank, the new translation
2 entry is created so that it targets physical memory from a previously untargeted
3 device or rank, and this device or rank may be restored to a non-reduced power
4 consumption mode.

5 In response to receiving a notification of a page de-allocation, memory
6 controller 14 deletes the corresponding mapping from table 30. Furthermore,
7 memory controller 14 is configured to periodically re-evaluate its existing memory
8 mappings to ensure that they are targeting the fewest possible number of physical
9 memory devices. The mappings are changed if necessary to achieve this goal,
10 after appropriately copying or moving affected portions of physical memory so
11 that each logical address will continue to reference the same data.

12 Re-evaluation of current memory mappings can be triggered in different
13 ways. For example, the memory controller might be configured to re-evaluate its
14 logical-to-physical address mappings at periodic intervals. Alternatively, the
15 mappings might be re-evaluated after each memory allocation or de-allocation,
16 after a pre-defined number of memory pages have been de-allocated, after a pre-
17 defined number of memory references, or whenever the power consumption of the
18 memory system approaches one or more pre-defined or dynamically determined
19 thresholds.

20 After adding or deleting mappings, or after re-evaluating the mappings as
21 described above, the memory controller identifies any memory devices or ranks of
22 memory devices that are not currently targeted by logical-to-physical memory
23 mappings. The memory controller then sets these memory devices or device ranks
24 to a reduced power consumption mode, such as a nap, standby, or power-down
25 mode. Alternatively, it might be desirable in some embodiments to identify those

1 devices or device ranks that are targeted by relatively few logical-to-physical
2 memory mappings, or by infrequent or non-recent memory accesses, and to set
3 those devices or device ranks to a reduced power mode. Although in these
4 examples the memory controller is responsible for identifying, tracking, and
5 placing devices into a reduced power mode, it is possible for software (such as the
6 operating system) to perform these functions as well.

7 8 **Use Registers to Indicate In-Use Memory**

9 In yet another embodiment, each memory device includes multiple
10 dynamically changeable use registers such as described in a co-pending US patent
11 application entitled "Monitoring In-Use Memory Areas for Power Conservation"
12 by inventors Steven C. Woo and Pradeep Batra, filed concurrently herewith, which
13 is hereby incorporated by reference.

14 Fig. 3 shows a memory device 40 that incorporates use registers 41 such as
15 those referred to in the patent application mentioned above. These registers
16 indicate used and unused memory cells or groups of used or unused memory cells.
17 More specifically, use registers 41 in this embodiment comprise individual bits or
18 flags that are associated respectively with individual memory cell rows 42. Each
19 bit or flag is set to indicate whether or not the corresponding row is actually in use,
20 and whether it therefore needs to be refreshed.

21 Memory controller 14 supports and maintains use registers 41 and allows
22 them to be set or programmed by the operating system to indicate which memory
23 rows are actually in use. For example, the operating system might set the use
24 registers to indicate which rows of logical memory are currently allocated. As
25 described in the previously mentioned patent application, use registers 41 allow

1 power-saving measures to be taken with respect to individual memory rows that
2 are not being used. Specifically, the use registers allow refreshing of unused
3 memory rows to be omitted.

4 Furthermore, in the described embodiment memory controller 14 monitors
5 use registers 41 to determine which portions of physical memory are in use at any
6 particular time, and periodically modifies logical-to-physical translation table 30
7 to consolidate in-use memory rows on as few memory devices or memory ranks as
8 possible. Remaining memory devices are then set to reduced power consumption
9 modes such as standby, nap, or power-down modes. Such reduced power modes
10 typically affect entire memory devices or ranks of memory devices, rather than
11 individual rows within such memory devices. However, it is possible that for
12 some designs, reduced power modes may affect portions of memory storage
13 within one or more devices, such as a bank within a DRAM.

14 In this embodiment, the address translation logic 22 initially implements
15 default logical-to-physical address mappings in table 30 of Fig. 2. Subsequently,
16 memory controller 14 periodically identifies or re-identifies portions of the logical
17 address space that are in use, by examining use registers 41. Furthermore,
18 memory controller 14 periodically re-maps the identified or re-identified portions
19 of the logical address space that are in use, in a manner that reduces the number of
20 physical memory devices or ranks that are referenced by the in-use portions of the
21 logical address space. Prior to re-mapping any particular range of logical
22 addresses, the memory controller moves any memory content that was previously
23 referenced by such logical addresses so that said memory content will continue to
24 be referenced by the same logical addresses after re-mapping.

1 Although the use registers or flags 41 are shown on an individual memory
2 device for purposes of this example, the registers or flags could be physically
3 located elsewhere. For example, they could be located on the memory controller
4 itself, or on some other component other than the memory controller or memory
5 device. Furthermore, such registers could be made to correspond to ranges of
6 logical memory addresses, rather than to physical address rows as shown in Fig. 3.

7 8 **Virtual Memory Implementation**

9 Fig. 4 shows another embodiment, in which similar techniques are
10 implemented within a virtual memory system of computer 10. Shown in Fig. 4 is
11 virtual memory logic 50 for implementing a virtual memory system. In systems
12 such as this, the operating system allocates virtual memory to requesting
13 application programs. Virtual memory can reside at times on secondary storage
14 media such as hard disks. Prior to use, however, the operating system loads
15 needed portions of the virtual memory into portions of physical memory, and
16 provides address translations from virtual memory addresses to physical memory
17 addresses. In this embodiment, memory controller 14 does not perform address
18 translations. Thus, CPU 12 specifies actual "physical" addresses to memory
19 controller 14, and the virtual memory system maintains mappings or translations
20 from virtual memory addresses to physical memory addresses. In this
21 embodiment, the virtual address space can be considered to be a "logical" address
22 space, and similar techniques are used to map from this "logical" or "virtual"
23 address space to the physical address space occupied by memory devices 16.

24 It should be noted that although the virtual memory logic is shown as being
25 implemented within operating system 18, the CPU 12 (and potentially other

1 components) typically also has logic for supporting the implementation of virtual
2 memory. Such CPU logic is typically under the control and supervision of the
3 operating system.

4 Fig. 5 shows a virtual-to-physical address translation table 60 that forms
5 part of virtual memory logic 50. This table has a plurality of mappings or entries
6 61, each of which indicates a virtual page address or number in a first column 62
7 and the corresponding targeted physical page address or number in a second
8 column 63. When referencing a given virtual address or virtual page number,
9 virtual memory logic 50 translates it to a physical address or physical page number
10 by referencing the appropriate entry of the address translation table 60. Virtual
11 memory logic 50 creates a mapping when a corresponding portion or page of
12 virtual memory is loaded into physical memory. When the portion or page of
13 virtual memory is removed from physical memory, the mapping is deleted from
14 table 60. Thus, address translation table 60 is a dynamically changing data
15 structure. At any given time, table 60 reflects only those virtual memory pages
16 that are actually loaded in physical memory. This changes as memory is allocated
17 and de-allocated, and as memory is moved back and forth between secondary
18 storage and physical memory. For purposes of the following discussion, physical
19 memory is considered to be allocated when it is the target of a virtual-to-physical
20 address mapping.

21 In one implementation, the virtual memory logic 50 of Fig. 4 is configured
22 to periodically and/or repeatedly identify portions of the virtual address space that
23 are in use, and to repeatedly modify the mappings of table 60 to reduce the number
24 of physical memory devices that are targeted by the mappings. This is similar to
25

1 the previously discussed embodiments, except that these steps are performed in the
2 virtual memory system rather than in the memory controller.

3 In another implementation, the mappings are created in a manner that tends
4 to minimize the number of in-use physical memory devices. In this
5 implementation, the virtual memory logic 50 manages its free page list so that
6 newly allocated pages are allocated from those memory devices that are already
7 receiving the heaviest usage.

8 Fig. 6 shows a free page or free region list 70 that is maintained by virtual
9 memory logic 50 in accordance with this implementation. This list contains
10 indications of free memory regions or pages—pages that are not currently targeted
11 by any mapping entry in table 60.

12 Free page list 70 is implemented as a linked list of free page or free region
13 indicators 72. Each free page indicator is a data structure containing a physical
14 address or page number and a pointer to the next indicator in the list (unless the
15 indicator is the last in the sequence).

16 The free page indicators are grouped, with each group corresponding to a
17 set of one or more physical memory devices. A set of memory devices typically
18 comprises those memory devices whose bits are combined to create a data word of
19 the width supported by the particular memory system architecture, and which
20 would therefore be set as a group to a particular power consumption mode. For
21 example, a set of memory devices might comprise a rank of memory devices.
22 Each group includes free memory indicators corresponding to free memory
23 regions or pages within the corresponding set of physical memory devices.

24 Fig. 6 shows four groups, corresponding to four different sets or ranks of
25 memory devices, Rank A, Rank B, Rank C, and Rank D, each of which is shown

1 in its own row. Within Rank A, there are two free physical memory pages. The
2 free page indicators corresponding to these two free memory pages are grouped
3 within the linked list, meaning that they are adjacent each other in the linked order.
4 Similarly, Rank B has four free pages, which are grouped or located adjacent in
5 the linked list.

6 Within the linked list, the indicators are sorted by group, in an order that is
7 based on the relative current memory allocations from the corresponding sets of
8 physical memory devices. More specifically, the groups are sorted in order from
9 the group corresponding to the memory rank having the most memory allocations
10 to the group corresponding to the memory rank having the fewest memory
11 allocations. In some cases, this is equivalent to an order beginning with the group
12 having the fewest free pages to the group having the most free pages.

13 In the example of Fig. 6, the first group in the sorted order corresponds to
14 the memory rank having the fewest number of free pages, which is Rank C. The
15 second group corresponds to the memory rank having the second fewest number
16 of free pages, which is Rank A. The last group in the order is that group having
17 the highest number of free pages, which in this example are Rank D and Rank B.

18 When allocating physical memory, the memory is allocated in the order of
19 the sorted free region list, so that physical memory is allocated preferentially from
20 those sets of physical memory devices having relatively higher current allocations
21 of memory regions. In the example of Fig. 6, memory is allocated first from Rank
22 C, then from Rank A, then from Rank D, and then from Rank B if necessary.

23 The sorting is preferably repeated periodically, as new physical memory
24 pages become available.
25

1 Allocating memory in this manner tends to consolidate memory usage on
2 relatively fewer physical memory devices. Periodically, the operating system or
3 virtual memory system identifies those ranks or devices having relatively fewer
4 memory allocations—those with the largest number of free pages, at the bottom of
5 the free region list—and sets those devices to reduced power modes. In some
6 cases, it might be desirable to identify only those memory devices or ranks that
7 have no current allocations, and to set only those devices to the reduced power
8 mode. In other cases, it is desirable to identify those memory devices or ranks that
9 have the least frequently and/or least recently accessed memory locations, and set
10 these devices to a reduced power mode.

11 In order to implement the described free page management technique, the
12 operating system will typically be configured with information regarding the
13 physical layout of memory, such as the address boundaries of memory ranks. This
14 information is used to determine the different groups of free page indicators.
15 Information regarding physical memory layout can be provided in different ways.
16 For example, the memory controller might be designed to support queries
17 regarding memory layout. In this case, the operating system would simply query
18 the memory controller for this information. Alternatively, memory layout
19 information might be provided by an operator as initial set-up information when
20 configuring a computer. As yet another alternative, memory layout information
21 may be provided as part of the system BIOS.

22 Although the technique of free page management has been described in the
23 context of an operating system and its virtual memory subsystem, similar free
24 page management techniques can be applied in the embodiments described with
25 reference to Figs. 1-3, in which a hardware-based memory controller creates

1 mappings from a logical address space to actual physical memory devices.
2 Specifically, a memory controller can maintain a sorted free page list
3 corresponding to free pages of physical memory. When creating new logical-to-
4 physical address mappings, the controller can create the mappings to first target
5 those physical memory pages at the top of the sorted list, to ensure that physical
6 memory is used first from those devices that are already heavily utilized.

7 Furthermore, it is recognized that the described techniques can in many
8 cases be implemented alternatively in software, hardware, or a combination of
9 both. In the case of software, an implementation might comprise one or more
10 computer-readable storage media containing instructions and instruction
11 sequences that are executable by a processor to perform the actions and techniques
12 described above. In the case of hardware, the same or similar actions might be
13 implemented as non-instruction-based logic components embodied on a hardware
14 substrate such as a circuit board or silicon chip.

15 16 **Conclusion**

17 The techniques described above can be used in many systems to produce
18 significant power savings. Furthermore, such power savings will often have few
19 or no detrimental side-effects, because the power-saving measures are taken with
20 respect to memory areas that are not actually being used or that are being used
21 relatively lightly. The described techniques can therefore avoid or greatly
22 diminish the prior art tradeoff between access speed and power savings.

23 Although details of specific implementations and embodiments are
24 described above, such details are intended to satisfy statutory disclosure
25 obligations rather than to limit the scope of the following claims. Thus, the

1 invention as defined by the claims is not limited to the specific features described
2 above. Rather, the invention is claimed in any of its forms or modifications that
3 fall within the proper scope of the appended claims, appropriately interpreted in
4 accordance with the doctrine of equivalents.
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25